

T10/04-136r0

SCSI to ATA Command Translations

April 27, 2004

T10 SCSI to ATA Translations
Study Group

Contact Information
Steve Johnson
LSI Logic Inc.
sjohnson@lsil.com
719 533 7511

SCSI to ATA Translations Study Group Meeting May 4

- What should this group produce?
 - Technical Paper or some type of standard moving forward?
 - Do we need an official project?
 - If a standard where does it belong?
- What version of specs should we translate to and from?
- What should this document cover?
- What should the format of the document be?
- Does anyone have interest in being the editor?

TABLE OF CONTENTS

1.	OVERVIEW	8
1.1	Notes	8
2.	SUMMARY OF COMMAND MAPPINGS	9
3.	COMMAND, RESULTS AND TRANSLATIONS	11
3.1	Inquiry command (12h)	11
3.1.1	Command Summary	11
3.1.2	SCSI INQUIRY CDB format and supported fields	11
3.1.3	Data Returned	11
3.1.3.1	Standard Inquiry data	11
3.1.3.2	Vital Product data – Supported VPD pages	13
3.1.3.3	Vital Product data – Unit serial number VPD page	14
3.1.4	Miscellaneous Notes	14
3.2	Read Capacity Command (10) and (25h)	15
3.2.1	Command Summary	15
3.2.2	Read Capacity (10) CDB format and supported fields	15
3.2.3	Data Returned	15
3.2.4	Miscellaneous Notes	15
3.3	Test Unit Ready Command (00h)	16
3.3.1	Command Summary	16
3.3.2	Test Unit Ready CDB format and supported fields	16
3.3.3	Data Returned	16
3.3.4	Miscellaneous Notes	16
3.4	Start Stop Unit Command (1Bh)	17
3.4.1	Command Summary	17
3.4.2	Start/Stop unit CDB format and supported fields	17
3.4.3	Data Returned	17
3.4.4	Miscellaneous Notes	17
3.5	Synchronize Cache Command (10) and (35h)	19
3.5.1	Command Summary	19
3.5.2	Synchronize Cache CDB format and supported fields	19
3.5.3	Data Returned	19
3.5.4	Miscellaneous Notes	19
3.6	Request Sense Command (03h)	20
3.6.1	Command Summary	20
3.6.2	Request Sense CDB format and supported fields	20
3.6.3	Data Returned	20
3.6.4	Miscellaneous Notes	20
3.7	Report LUNS Command (A0h)	21
3.7.1	Command Summary	21
3.7.2	Report LUNs CDB format and supported fields	21
3.7.3	Data Returned	21
3.7.4	Miscellaneous Notes	21
3.8	Format Unit Command (04h)	22
3.8.1	Command Summary	22
3.8.2	Format Unit CDB format and supported fields	22
3.8.3	Data Returned	22
3.8.4	Miscellaneous Notes	22
3.9	Verify Command (10) and (2Fh)	24
3.9.1	Command Summary	24
3.9.2	Verify CDB format and supported fields	24
3.9.3	Data Returned	24
3.9.4	Miscellaneous Notes	24
3.10	Read Command (6, 10, 12) (08h, 28h, A8h)	26
3.10.1	Command Summary	26
3.10.2	Read CDB format and supported fields	26
3.10.3	Data Returned	27
3.10.4	Miscellaneous Notes	27
3.11	Write (06, 10, 12) Command (0Ah, 2Ah, AAh)	28
3.11.1	Command Summary	28
3.11.2	Write CDB format and supported fields	28
3.11.3	Data Returned	28
3.11.4	Miscellaneous Notes	29
3.12	Write Verify Command (2Eh)	30

3.12.1	Command Summary	30
3.12.2	Write CDB format and supported fields.....	30
3.12.3	Data Returned	30
3.12.4	Miscellaneous Notes.....	30
3.13	Read Buffer (10) Command (3Ch).....	31
3.13.1	Command Summary	31
3.13.2	Read Buffer CDB format and supported fields.....	31
3.13.3	Data Returned	32
3.13.3.1	Data Only Mode (02h)	32
3.13.3.2	Descriptor mode (03h)	32
3.13.4	Miscellaneous Notes.....	32
3.14	Write Buffer (10) Command (3Bh).....	33
3.14.1	Command Summary	33
3.14.2	Read Buffer CDB format and supported fields.....	33
3.14.3	Data written.	33
3.14.3.1	Data Only Mode (02h)	34
3.14.4	Miscellaneous Notes.....	34
3.15	Seek (06, 10) Command (0Bh, 2Bh)	35
3.15.1	Command Summary	35
3.15.2	Seek CDB format and supported fields	35
3.15.3	Data Returned	35
3.15.4	Miscellaneous Notes.....	35
3.16	Rezero Unit (06) Command (01h).....	36
3.16.1	Command Summary	36
3.16.2	Rezero CDB format and supported fields.....	36
3.16.3	Data Returned	36
3.16.4	Miscellaneous Notes.....	36
3.17	Reserve Unit Command (16h).....	37
3.17.1	Command Summary	37
3.17.2	Reserve CDB format and supported fields.....	37
3.17.3	Data Returned	37
3.17.4	Miscellaneous Notes.....	37
3.17.5	SAS Affiliation Consideration.....	37
3.18	Release Unit Command (17h) (06/10)	38
3.18.1	Command Summary	38
3.18.2	Reserve CDB format and supported fields.....	38
3.18.3	Data Returned	38
3.18.4	Miscellaneous Notes.....	38
3.18.5	SAS Affiliation Considerations	38
3.19	Send Diagnostic Command (1Dh).....	39
3.19.1	Command Summary	39
3.19.2	Send Diagnostic CDB format and supported fields.....	39
3.19.3	Data Returned	40
3.19.4	Miscellaneous Notes.....	40
3.20	Log Sense Command (4Dh).....	41
3.20.1	Command Summary	41
3.20.2	Log Sense CDB format and supported fields	41
3.20.3	Data Returned	42
3.20.3.1	SMART data.....	42
3.20.4	Miscellaneous Notes.....	42
3.21	Mode Sense (06)/(10) Command (1Ah)/(5Ah).....	43
3.21.1	Command Summary	43
3.21.1.1	Mode Sense (06) CDB format and supported fields	43
3.21.1.2	Mode Sense (10) CDB format and supported fields	43
3.21.2	Data Returned	44
3.21.2.1	Reply for Mode Sense (06) Command	44
3.21.3	Miscellaneous Notes.....	45
3.22	Mode Select (06)/(10) Command (15h)/(55h).....	46
3.22.1	Command Summary	46
3.22.1.1	Mode Select (06) CDB format and supported fields	46
3.22.1.2	Mode Select (10) CDB format and supported fields	46
3.22.2	Data Transferred	46
3.22.2.1	Format for Mode Select (06) Command.....	46
3.22.3	Miscellaneous Notes.....	47
4.	SUPPORTED SCSI MODE PAGES	49
4.1	Commonly used SCSI Mode Pages	49

4.2	Read-Write Error Recovery Mode Page (01h)	50
4.2.1	Read-Write Error Recovery Mode Page format and supported fields	50
4.3	Format Mode Page (03h)	51
4.4	Rigid Disk Device Geometry Page (04h)	52
4.5	Caching Mode Page (08h)	52
4.6	Informational Exceptions Control Mode Page (1Ch)	54
5.	SCSI TASK MANAGEMENT TRANSLATION	56
5.1	Abort Task	56
5.2	Abort Task Set	56
5.3	Logical Unit Reset	56
5.4	Target Reset	56
5.5	Bus Reset	56
5.6	Clear Task Set	56
6.	SMART HANDLING	58
6.1	General Description of SMART	58
6.2	Host Notification	58
6.3	Disabling, enabling SMART	58
6.4	Retrieving SMART data from targets	58

LIST OF TABLES

Table 3.1.1.1 INQUIRY Command summary.....	11
Table 3.2.1.2 Read Capacity (10) Command summary.....	15
Table 3.3.1.3 Test Unit ready command summary.....	16
Table 3.4.1.4 Start/stop unit command summary.....	17
Table 3.4.2.5 Interpretation of START and LOEJ bit in start/stop unit CDB.....	17
Table 3.5.1.6 Synchronize Cache command summary	19
Table 3.6.1.7 Request sense command summary.....	20
Table 3.7.1.8 Report LUNs command summary	21
Table 3.8.1.9 Format unit command summary.	22
Table 3.9.1.10 Verify command summary.....	24
Table 3.10.1.11 Read command summary.....	26
Table 3.11.1.12 Write command summary	28
Table 3.12.1.13 Write Verify command summary	30
Table 7.13.1.14 Read Buffer command summary	31
Table 7.14.1.15 Write Buffer command summary.....	33
Table 3.13.1.14 Seek command summary.....	35
Table 3.14.1.15 Rezero command summary	36
Table 3.15.1.16 Reserve Unit command summary.....	37
Table 3.16.1.17 Reserve Unit/ Release Unit command summary	38
Table 3.17.1.18 Send diagnostic command summary.....	39
Table 3.17.2.19 Send diagnostic self-test code.....	39
Table 3.18.1.20 Log Sense command summary.....	41
Table 3.18.2.21 Page Control Value	41
Table 3.18.2.22 Page Code Values	41
Table 3.19.1.23 Mode Sense (06)/(10) commands summary	43
Table 3.19.1.24 Page Control (PC) values and their descriptions	43
Table 3.20.1.25 Mode Select (06)/(10) commands summary.....	46
Table 4.1.1 Commonly used SCSI mode pages	49

1. OVERVIEW

This document is intended to provide a working set of SCSI commands to ATA command translations for the purpose of emulating ATA devices as SCSI devices. This is becoming a common design choice when considering Serial Attached SCSI (SAS) topologies ability to support both command sets using SSP and STP protocols. Emulating a ATA as SCSI allows for SAS and SATA solutions to use a common command set in the upper software layers and host bus adapters (HBA) interfaces. This allows for solutions to preserve the current OS tried and true SCSI driver stacks and supporting management software.

There are many existing solutions that use SCSI to ATA command translations. These solutions have been designed independently of each other and typically vary only slightly in the command translations. This document is intended as a starting point for the T10 SCSI to ATA translation study group. The T10 study group was formed to explore the need and possibilities of the storage industry coming to a consensus on a single consistent programming model for future design to be based from and validated against. This will provide known and expected behavior to design to and test against.

Special thanks to Samantha Ranaweera of LSI Logic for crafting the original document this document leveraged so heavily from.

1.1 Notes

These notes apply to all command translations described in this document unless otherwise noted.

- In the event of a discrepancy between the contents of this section and the description of individual commands, description of individual commands takes precedence.
- This document describes the SCSI to SATA command mapping. No reference should be made to the handling of ATAPI devices from this document.
- Only LUN 0 is supported for all SATA targets. If a command tries to access a LUN other than 0, a check condition with the sense key set to illegal request and additional sense code set to invalid field in CDB is returned.
- Unless otherwise noted, the IMMED bit (immediate return) is ignored. For the format unit command, this bit will be supported.
- All ATA commands with the exception of SATA Native Command Queued Reads and Writes are single threaded per device. The translator will queue IOs as necessary to enforce this.

2. SUMMARY OF COMMAND MAPPINGS

Table 1.1 Summary of SCSI-SATA command mapping

SCSI COMMAND	SCSI Opcode	ATA Opcode	Support Mechanism		Hyper Link
			Emulate	Translate	
Change Definition	40h	N/A			Not Supported
Compare	39h	N/A			Not Supported
Copy	18h	N/A			Not Supported
Copy and Verify	3Ah	N/A			Not Supported
Format Unit	04h	35h, 36h	X	X	3.8
Inquiry	12h	ECh	X		3.1
Lock-unlock cache	36h	N/A			Not Supported
Log Select	4Ch	N/A			Not Supported
Log Sense	4Dh	N/A	X		3.20
Mode Select (6,10)	15h, 55h	N/A	X		3.22
Mode Sense (6,10)	1Ah, 5Ah	N/A	X		3.21
Move Medium	A7h	N/A			Not Supported
Persistent Reserve In	5Eh	N/A			Not Supported
Persistent Reserve Out	5Fh	N/A			Not Supported
Prefetch	34h	N/A			Not Supported
Prevent Allow Medium Removal	1Eh	N/A			Not Supported
Read (6,10,12)	08h, 28h, A8h			X	3.10
Read Buffer	3Bh	E4h	X	X	3.13
Write Buffer	3Ch	E8h	X	X	3.14
Read Capacity	25h	N/A	X		3.2
Read Defect Data (10)	37h	N/A			Not Supported
Read Defect Data (12)	B7h	N/A			Not Supported
Read Element Status	B4h	N/A			Not Supported
Read Long (10)	3Eh	N/A			Not Supported
Reassign blocks	07h	N/A			Not Supported
Rebuild	81h	N/A			Not Supported
Receive diagnostic results	1Ch	N/A			Not Supported
Regenerate	82h	N/A			Not Supported
Release and Reserve (6,10)	16h, 17h, 56h, 57h	N/A			Not Supported
Report Device Identifier	A3h	N/A			Not Supported
Report LUNs	A0h	N/A	X		3.7

Request Sense	03h	N/A	X		3.6
Rezero Unit	01h	N/A	X		3.16
Search Data Equal	31h	N/A			Not Supported
Search Data High	30h	N/A			Not Supported
Search Data Low	32h	N/A			Not Supported
Seek (6,10)	0Bh, 2Bh	N/A	X		3.15
Send Diagnostic	1Dh	90h		X	3.19
Set Device Identifier	A4h	N/A			Not Supported
Set Limits	33h	N/A			Not Supported
Start Stop Unit	1Bh	E0h, E1h		X	3.4
Synchronize Cache	35h	E7h		X	3.5
Test Unit Ready	00h	E5h		X	3.3
Verify (10)	2Fh	40h, 42h		X	3.9
Verify (12,16)	AFh, 8Fh	N/A			Not Supported
Write (6,10,12)	0Ah, 2Ah, AAh	Multiple		X	3.11
Write and Verify (10)	2Eh	Multiple		X	3.12
Write and Verify (12,16)	A Eh, 8 Eh	N/A			Not Supported
Write Long	3Fh	N/A			Not Supported
Write Same (10,12)	41h, 93h	N/A			Not Supported
XdRead	52h	N/A			Not Supported
XdWrite	50h	N/A			Not Supported
XdWrite extended	80h	N/A			Not Supported
XpWrite	51h	N/A			Not Supported

3. COMMAND, RESULTS AND TRANSLATIONS

3.1 Inquiry command (12h)

The SCSI Inquiry command requests general information about a target or component LUN. Inquiry and certain vital product data pages have translations.

3.1.1 Command Summary

Table 3.1.1.1 INQUIRY Command summary

Support Method	Emulated
ATA Opcodes	ECh, ATA Identify Device Command.
Supported Vital Product Data pages	00h, Supported VPD pages 80h, Unit serial number

3.1.2 SCSI INQUIRY CDB format and supported fields

- CmdDt

This bit is now obsolete and is not supported. If set to one, return a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

- EVPD

This bit must be set to request vital product data pages and the Page Code should indicate the page requested. Only two pages are supported currently: Supported VPD pages (00h) and Unit serial number (80h).

3.1.3 Data Returned

Standard inquiry data or vital product pages will be returned depending on the request.

3.1.3.1 Standard Inquiry data

- PERIPHERAL DEVICE TYPE

This bit field is set to 0 to indicate that a direct access device is connected.

- PERIPHERAL QUALIFIER

This bit field is set to 0 to indicate that the specified peripheral device is currently attached to this logical unit.

- RMB

Removable medium bit. Set the value of bit 7 of General Configuration field of IDENTIFY DEVICE information retrieved from the target.

- VERSION

This field is set to 03h to indicate that it supports SCSI 3 SPC specification.

- RESPONSE DATA FORMAT

Set to a value of 2. Values other than 2 are obsolete. Value of 2 indicates that data shall be in the format supported by SPC-3.

- HiSup (hierarchical support)

This bit is set to 0 to indicate that this device shall not support hierarchical addressing model.

- NORMAL ACA SUPPORT

This bit is set to 0 to indicate the NACA bit in the control byte of the SCSI inquiry CDB is not supported.

- ADDITIONAL LENGTH

This field is set to indicate the length of the inquiry data that follows.

- PROTECT

Set to zero to indicate that the device does not support protection information.

- 3PC (Third party commands)

Set to 0 to indicate that this device does not support third party commands.

- ALUA

These bits are set to 0 to indicate that this device does not support asymmetrical logical unit access.

- ACC

This bit is set to 0 to indicate that no access controls coordinator may be addressed through this logical unit.

- SCC

This bit is set to 0 to indicate that the device does not contain an embedded storage array controller component.

- Add16

This bit is set to 0 to indicate that this device does not support 16bit wide SCSI addresses.

- MCHNGR

This bit is set to 0 to indicate that this device is not attached to medium transport element.

- MULTIP

This bit is set to 0 to indicate that this device does not contain multiple ports.

- ENCSERV

This bit is set to 0 to indicate that this device does not contain an embedded enclosure service component.

- BQUE

Works in combination with the CMDQUE bit. This bit is set to 0 to indicate that this device does not support basic task management model.

- CMDQUE

This bit is set if bit 8 of the SerialATACapabilities field in the device Identify data is set and the SerialATACapabilities field represents a valid value. If set to 1 the device supports full task management model, specifically it supports command queuing. If this bit is 0, the device does not support command queuing.

- LINKED

This bit is set to 0 to indicate that this device does not support linked commands.

- SUNC

This bit is set to 0 to indicate that this device does not support SCSI synchronous data transfers.

- WBUS16

This bit is set to 0 to indicate that this device does not support 16 bit transfers. SYNC/WBUS16 applies to parallel SCSI.

- VENDOR ID

This 8 bytes field is created using the first 8 bytes of the 40-byte Model Number field from the device identify data retrieved from the device. Endianness is swapped to match the SCSI format.

- PRODUCT ID

This 16 bytes field is created using the bytes 8 through 23 of the 40 byte Model Number field from the identify data. Endianness is swapped to match the SCSI format.

- PRODUCT REVISION

This field is created using the 8-byte firmware revision field of the device Identify data. The lower 4 bytes are used and the upper 4 bytes are lost since SCSI supports only 4 bytes for revision number. Endianness is swapped to match the SCSI format.

- IUS/QAS/CLOCKING/Vendor descriptors/ and vendor specific parameters of standard inquiry data are not supported.

3.1.3.2 Vital Product data – Supported VPD pages

- PERIPHERAL DEVICE TYPE

This bit field is set to 0 to indicate that a direct access device is connected.

- PERIPHERAL QUALIFIER

This bit field is set to 0 to indicate that the specified peripheral device is currently attached to this logical unit.

- PAGE CODE

This field is set to 0 to indicate '*supported VPD*' pages page.

- PAGE LENGTH

This field indicates the length of the supported VPD page list returned in number of bytes.

- SUPPORTED VPD LIST

This list contains the page codes of the pages supported and is implemented in ascending order of page codes beginning with page code 00h.

3.1.3.3 Vital Product data – Unit serial number VPD page

- PERIPHERAL DEVICE TYPE

This bit field is set to 0 to indicate that a direct access device is connected.

- PERIPHERAL QUALIFIER

This bit field is set to 0 to indicate that the specified peripheral device is currently attached to this logical unit.

- PAGE CODE

This field is set to 80h to indicate '*supported VPD*' pages page.

- PAGE LENGTH

This field indicates the length of the product serial number. Currently a 20-byte product serial number is being returned; hence the length is set to 20 bytes.

- PRODUCT SERIAL NUMBER

The 20-byte field, serial number, found in the device identify data retrieved from the device will be returned in this field. The serial number field will be byte swapped to match the SCSI data format.

3.1.4 Miscellaneous Notes.

The LUN should be set to zero. If set, return SCSI_INQUIRY_PERIPHERAL_QUALIFIER_NOT_SUPPORTED and SCSI_INQUIRY_DEV_TYPE_UNKNOWN (7Fh).

3.2 Read Capacity (10) Command (25h)

Read Capacity (10) command requests information about the capacity of the block device being addressed. Information protection is not supported therefore the read capacity (16) command is not supported.

3.2.1 Command Summary

Table 3.2.1.2 Read Capacity (10) Command summary

Support Method	Emulated
ATA Opcodes	None. Use ATA Identify Device Information at power-on or insertion to compute the maximum user addressable medium capacity.

3.2.2 Read Capacity (10) CDB format and supported fields

- LOGICAL BLOCK ADDRESS

This field works in conjunction with the PMI bit described below. Since this requires vendor specific information this operation is not supported and shall be set to zero.

- PMI (Partial Medium Indicator)

Operation specified by this bit is ignored. (Return data for the LBA, after the LBA specified in LOGICAL BLOCK ADDRESS discussed above, for which a vendor specific delay in data transfer may be encountered)

3.2.3 Data Returned

Read Capacity data as defined by SCSI Block Commands specification will be returned to the requester.

- RETURNED LOGICAL BLOCK ADDRESS

TotalUserSectors-1 (words 60/61) found in the device Identify data retrieved from the device, Endianness swapped to match SCSI format.

- BLOCK LENGTH IN BYTES

This value is currently set to 512 bytes, which is the standard sector size for disk drives.

3.2.4 Miscellaneous Notes.

Ignore PMI bit and LOGICAL BLOCK ADDRESS in the incoming SCSI CDB. Read Capacity (16) command is not supported.

3.3 Test Unit Ready Command (00h)

Test unit command is used to determine whether the device is ready or not. If the device is capable of accepting medium access commands without failing them with a check condition then the command should return GOOD status. However, if the device is unable to process medium access commands, Test Unit Ready Command should return a check condition with sense key set to NOT READY.

3.3.1 Command Summary

Table 3.3.1.3 Test Unit ready command summary

Support Method	Translated to ATA Check Power mode Command.
ATA Opcodes	E5h

3.3.2 Test Unit Ready CDB format and supported fields

All fields in the CDB shall be set to zero.

3.3.3 Data Returned

If the device is capable of accepting medium access commands without failing with a check condition, then a status of GOOD shall be returned.

If the device is not capable of accepting medium access commands, then check condition shall be returned with the sense key set to not ready and additional sense code set to logical unit not ready.

3.3.4 Miscellaneous Notes.

- Test unit ready command is mapped to the ATA Check Power Mode command. If the drive is at standby power mode, then it is assumed that the device is incapable of accepting medium access commands without any intervention, so a check condition is returned with the sense key set to not ready and the additional sense code set to logical unit not ready, initializing command required.
- If the ATA Check power mode command returned a power state of active or idle, then it is assumed that the drive is ready to accept medium access commands, so a status of good shall be returned to the requester.
- If the device is being formatted, a check condition is returned with the sense key set to not ready and the additional sense code set to logical unit not ready, format in progress.

3.4 Start Stop Unit Command (1Bh)

The start stop unit provides a method for controlling the power state of a logical unit. For SATA devices, this implies transition between idle and standby power modes. This command is also used to load and unload medium as well. Load medium request will return a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

3.4.1 Command Summary

Table 3.4.1.4 Start/stop unit command summary

<u>Support Method:</u>	<u>Translated to:</u>
<ul style="list-style-type: none"> Start Unit Stop Unit Unload Media Load Media 	<ul style="list-style-type: none"> Idle Immediate Standby immediate Check condition with sense key illegal request and ASC invalid field in CDB Check condition with sense key illegal request and ASC invalid field in CDB
ATA Opcodes	E0h, E1h

3.4.2 Start/Stop unit CDB format and supported fields

- IMMED – Immediate bit
Ignored.
- START/LOEJ

Table 3.4.2.5 Interpretation of START and LOEJ bit in start/stop unit CDB

LOEJ	START	Interpretation
0	0	Stop Unit Command. Translated to standby immediate
0	1	Start Unit Command. Translated to idle Immediate.
1	0	Not supported. Check condition returned with sense key illegal request.
1	1	Not supported, Check condition returned with sense key illegal request.

- POWER CONDITIONS
Ignored.

3.4.3 Data Returned

This command does not return any data. Status or a check condition is returned after the completion of the command. The immediate bit in the CDB is ignored.

3.4.4 Miscellaneous Notes.

- POWER CONDITION
Ignored

3.5 Synchronize Cache (10) Command (35h)

Synchronize cache (10) command is used to flush the most recent data values in the device cache to physical medium. Unlike in SCSI, a particular LBA cannot be specified to start flushing the device cache.

3.5.1 Command Summary

Table 3.5.1.6 Synchronize Cache command summary

Support Method	Translated into flush cache command.
ATA Opcodes	E7h

3.5.2 Synchronize Cache CDB format and supported fields

- IMMED Immediate Bit
This bit is ignored.
- LOGICAL BLOCK ADDRESS
This field is ignored.
- NUMBER OF BLOCKS
This field is ignored.

3.5.3 Data Returned

This command does not return any data. Status is returned after the completion of the command.

3.5.4 Miscellaneous Notes.

LBA and NUMBER OF BLOCK fields are ignored, as they cannot be translated to ATA command fields.

3.6 Request Sense Command (03h)

Request sense command requests that any available sense data be returned to the caller. Upon errors with auto request sense enabled, the translator will issue a request sense command and send the sense data to the caller along with the status of the command that generated the check condition. This behavior is emulated for SATA devices as well. Therefore, a separate request sense command is not necessary and will result in a sense data block with sense key set to no sense and additional sense code set to no additional sense data. Emulate enabling and disabling of auto request sense and cache sense data if auto request sense is disabled.

The only exception to this is when a unit format command is being executed or when SMART reporting is enabled. It is legal to send a request sense command to a target device while a device format command is outstanding. If requested, return a sense key of not ready and additional sense code set to logical unit not ready format in progress. The sense key specific bytes will be set to progress indication as per SPC-3. Return sense data with sense key set to no sense, additional sense key set to general hardware failure if it detects a SMART threshold exceeding condition. Refer section 6 on SMART for further details.

3.6.1 Command Summary

Table 3.6.1.7 Request sense command summary

Support Method	Emulated. Auto request sense is always enabled, so a separate request sense command will result in sense data with sense key set to no sense. Only exception is when a format unit command is outstanding and when SMART reporting is enabled.
ATA Opcodes	None.

3.6.2 Request Sense CDB format and supported fields

- DESC

This bit is ignored.

- ALLOCATION LENGTH

This field specifies how many bytes have been allocated for the returned sense data. The minimum of this value of 18 bytes shall be returned.

3.6.3 Data Returned

Sense data will be returned with the sense key set to no sense and additional sense code set to no additional sense data. When a format unit command is outstanding, sense data indicating the progress of the format operation will be returned. If SMART reporting is enabled, and a threshold exceeding condition is detected, key set to no sense and additional sense code set to general hardware failure will be returned.

3.6.4 Miscellaneous Notes.

None.

3.7 Report LUNS Command (A0h)

Report LUNs command is used to discover the logical unit inventory of a peripheral device. Logical unit inventory is returned to the requester as a list of logical unit numbers for all logical units having a PERIPHERAL QUALIFIER of 000b. Logical units with other values for PERIPHERAL QUALIFIER are currently not included in this list.

3.7.1 Command Summary

Table 3.7.1.8 Report LUNs command summary

Support Method	Emulated. Only LUN 0 is supported currently.
ATA Opcodes	None.

3.7.2 Report LUNs CDB format and supported fields

- SELECT REPORT

This field is ignored.

- ALLOCATION LENGTH

This field should be at least 16 bytes according to the SPC-3 specification and that should guarantee that the length of the user buffer used to DMA the report LUNs data back is not overrun.

3.7.3 Data Returned

Data returned will be in compliance to the REPORT LUNs parameter data format specified in SPC-3. Currently the LUN LIST LENGTH, which is 8 times the number of logical units reported back, is set to 8 and only LUN 0 is reported back to the caller.

3.7.4 Miscellaneous Notes.

Only LUN zero is reported.

3.8 Format Unit Command (04h)

Format Unit Command verifies that all logical block addresses visible to external application clients are formatted and can be accessed. All sectors of the visible address space are written to zero.

3.8.1 Command Summary

Table 3.8.1.9 Format unit command summary.

Support Method	Emulated. It will use the same Write commands as in the section 3.11. Depending on if Extended and Queued bits are set, different ATA opcodes will be used. No optional parameters implemented at this time except for the IMMED bit.
ATA Opcodes	None. 35h used to zero out all sectors in the visible address space.

3.8.2 Format Unit CDB format and supported fields

- DEFECT LIST FORMAT

This feature is not supported. If requested terminate the command with a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

- CMLIST

This feature is not supported. If requested terminate the command with a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

- FMTDATA

If set to 0, no data shall be transferred from the data-out buffer. If set to 1, format unit parameter list will be transferred from the client's data out buffer. Note that only the IMMED bit in the parameter list will be supported. All other fields are ignored.

- LONGLIST

This field is currently not supported. If requested terminate the command with a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

- FMTPINFO

This feature is not supported. If requested terminate the command with a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB.

- INTERLEAVE

This field is ignored.

3.8.3 Data Returned

None.

3.8.4 Miscellaneous Notes.

During the process of formatting a device Inquiry commands are processed as normal. Request sense commands will be processed with Not Ready and Format in Process sense key and additional sense

code. All other commands will result in a check condition with the same sense key and additional sense key. Optional, to reply with a check condition a sense key of illegal request if an unsupported field or a bit is set to a value other than what is defined in this document.

3.9 Verify (6,10,12) Command (2Fh)

Verify command is used to verify data on medium, which includes user data and protection data. This SCSI command is directly translated into ATA read verify sectors command or the extended version of that command.

3.9.1 Command Summary

Table 3.9.1.10 Verify command summary

Support Method	Translated in ATA read verify sector(s) command. If the device supports 48 bit mode, translated into ATA read verify extended command.
ATA Opcodes	40h for read verify sector command. 42h for read verify Ext. command.

3.9.2 Verify CDB format and supported fields.

- DPO (Disable pageout)
Ignored.
- BLKVFY
Ignored.
- BYTCHK
Ignored.

3.9.3 Data Returned

None.

3.9.4 Miscellaneous Notes.

- If the LOGICAL BLOCK ADDRESS is larger than what can be accommodated using 28 bits and the device does not support 48-bit mode, a check condition with sense key set to illegal request and additional sense code set to logical block address out of range shall be returned.
- If the LBA plus the VERIFICATION LENGTH is greater than the maximum sector that can be addressed on medium, return a check condition with sense key set to illegal request and additional sense code set to logical block address out of range.
- LBA mapping algorithm assigns LSB of the LBA field from the SCSI CDB, i.e. byte 5, to LBA_{Low} in the ATA FIS, byte 4 from SCSI CDB to LBA_{Mid}, byte 3 from SCSI CDB to LBA_{High}. If the device supports 48-bit addressing, then byte 2 of the SCSI CDB is assigned to LBA_{LowExp}.
- SectorCount in the ATA FIS is derived from the LSB, i.e. byte 8, of the VERIFICATION LENGTH of the SCSI CDB. If the device supports 48-bit addressing, the SectorCountExp is assigned byte 7 of the SCSI CDB, which is the MSB of the VERIFICATION LENGTH field.
- Only Read/Write commands can be queued, so a performance degradation is expected when commands such as Verify is issued as they have to run in a single threaded manner.

3.10 Read (6, 10, 12) Command (08h, 28h, A8h)

The read command is used to request the device to transfer user data to the requester. Different versions of the command support different LBA sizes and different transfer lengths. Data can be read from medium or the device cache if the most recent copy is in the cache and has not been transferred to the medium.

3.10.1 Command Summary

Table 3.10.1.11 Read command summary

Support Method	Translated to ATA read commands using the following criterion.			
ATA Opcodes	<i>Is Queuing enabled</i>	<i>Are Extended Commands Enabled</i>	<i>SCSI OP Code</i>	<i>Translated ATA Opcode</i>
	NO	NO	08h,28h,A8h	READ_DMA (C8h)
	NO	YES	08h,28h,A8h	READ_DMA_EXT (25h)
	YES	(Don't care)	08h,28h,A8h	READ_FPDMA_QUEUED

3.10.2 Read CDB format and supported fields.

- DPO (Disable page out)
Ignored.
- FUA (Force Unit Access)
Supported with NCQ targets.
For non NCQ drives do we need to send a cache flush to force unit access?
- RDPROTECT
Ignored.
- TRANSFER LENGTH
 - For Read(06) Commands (08h), transfer length of 0 is interpreted as 256 sectors and is the same for SCSI and ATA.
 - For Read(10) Commands (28h), transfer length of 0 indicates that no data transfer should take place.
 - For Read(06) and Read(10) commands, if the transfer length is not zero, it should indicate the number of sectors to transfer from the device to the requester.
- LOGICAL BLOCK ADDRESS
 - For Read(06) Commands (08h), a 21 bit LBA is derived from bytes 3, 2, 1:5 in the CDB, with byte 3 being the LSB and the 5 bits from byte 1 being the MSB.
 - For Read(10) Commands (28h), A 32 bit LBA is derived from bytes 2 through 5, where byte 5 is the LSB and byte 2 is the MSB.

- For Read(10) Commands (28h), A 32 bit LBA is derived from bytes 2 through 5, where byte 5 is the LSB and byte 2 is the MSB.

3.10.3 Data Returned

Requested sector(s) will be returned if successfully retrieved from the medium. Check conditions will be reported back if the LBA or the LBA plus the TRANSFER LENGTH is greater than the maximum sector that can be addressed in medium with the sense key set to illegal request and additional sense code set to invalid field in CDB.

3.10.4 Miscellaneous Notes.

Protection information is not supported. For the Read(12) command, if the size TRANSFER LENGTH field is greater than 16 bits, then the command will terminate with a check condition where the sense key set to illegal request and additional sense code set to invalid field in CDB. This is because the ATA command set support a maximum of 16-bit sector count field.

3.11 Write (06, 10, 12) Command (0Ah, 2Ah, AAh)

The write command is used to request the device to transfer user data to device medium or cache. Different version of the command supports different LBA sizes and or different transfer lengths. Data can be written to medium or the device cache.

3.11.1 Command Summary

Table 3.11.1.12 Write command summary

Support Method	Translated to ATA write commands using the following criterion.			
ATA Opcodes	<i>Is Queuing enabled</i>	<i>Are Extended Commands Enabled</i>	<i>SCSI OP Code</i>	<i>Translated ATA Opcode</i>
	NO	NO	0Ah,2Ah,AAh	WRITE_DMA (C8h)
	NO	YES	0Ah,2Ah,AAh	WRITE_DMA_EXT (25h)
	YES	(Don't care)	0Ah,2Ah,AAh	WRITE_FPDMA_QUEUED

3.11.2 Write CDB format and supported fields.

- DPO (Disable page out)
Ignored.
- FUA (Force Unit Access)
Supported with NCQ capable devices.
- WRPROTECT
Ignored.
- TRANSFER LENGTH
 - For Write(06) Commands (0Ah), transfer length of 0 is interpreted as 256 sectors and is the same for SCSI and ATA.
 - For Write(10) Commands (2Ah), transfer length of 0 indicates that no data transfer should take place.
 - For Write(06) and Write(10) commands, if the transfer length is not zero, it should indicate the number of sectors to transfer from the device to the requestor.
- LOGICAL BLOCK ADDRESS
 - For Write(06) Commands (0Ah), a 21 bit LBA is derived from bytes 3, 2, 1:5 in the CDB, with byte 3 being the LSB and the 5 bits from byte 1 being the MSB.
 - For Read(10) Commands (2Ah), A 32 bit LBA is derived from bytes 2 through 5, where byte 5 is the LSB and byte 2 is the MSB.

3.11.3 Data Returned

Requested sector(s) will be transferred to the addressed device, and the device may transfer the data to its cache or medium. Check condition will be reported back if the LAB or the LBA plus the

TRANSFER LENGTH is greater than the maximum sector that can be addressed in medium, where the sense key set to illegal request and additional sense code set to invalid field in CDB.

3.11.4 Miscellaneous Notes.

- Note that protection information is not supported. Also note that SCSI forced unit access mechanism is not supported.
- Note that for Write (12) command, if the size TRANSFER LENGTH field is greater than 16 bits, then the command will terminate with a check condition where the sense key set to illegal request and additional sense code set to invalid field in CDB. This is because the ATA command set support a maximum of 16-bit sector count field.

3.12 Write Verify Command (2Eh)

The write verify command is used to transfer application data to medium and then to verify that data was written correctly.

3.12.1 Command Summary

Table 3.12.1.13 Write Verify command summary

Support Method	Translated to a ATA write command first and then to read verify sector(s) command.
ATA Opcodes	See section 3.11 for details on how the ATA write command is derived. 40h for read verify sector command if the device does not support extended commands feature set. 42h for read verify Ext. command if the device supports the extended Commands feature set.

3.12.2 Write CDB format and supported fields.

- BYCHK
Ignored.
- EBP
Ignored.
- DPO (Disable pageout)
Ignored.
- LOGICAL BLOCK ADDRESS
See section 3.11 for details on how this field is translated into ATA write command LBA fields.
- TRANSFER LENGTH
Translated into sector count of the ATA command. Note that TRANSFER LENGTH of 0 indicates that no data will be transferred.

3.12.3 Data Returned

Data will be transferred from the application client's buffer to medium. Only status will be returned to the application client.

3.12.4 Miscellaneous Notes.

Protection information is not supported.

3.13 Read Buffer Command (3Ch)

The read buffer command is used with the write buffer command to determine the integrity of the target device's buffer memory and the physical interconnect that connects the target device and the initiator.

3.13.1 Command Summery

Table 3.13.1.14 Read Buffer command summery

Support Method	Partially translated, partially emulated.		
ATA Opcodes	MODE	Supported Method	Translated ATA Opcode
	00h – Combined Header and Data Mode	Not Supported.	N/A
	01h – Vendor specific Mode	Not Supported.	N/A
	02h – Data Only mode	Translated to ATA Read Buffer Command	E4h
	03h – Descriptor mode	Emulated.	N/A
	0Ah – Echo Mode	Not supported.	N/A
	0Bh – Echo buffer descriptor Mode	Not Supported.	N/A
	All other MODE values	Not Supported.	N/A

3.13.2 Read Buffer CDB format and supported fields

- **MODE**

Please refer

Table 3.13.1.14 for a description of what MODEs are supported. Only data, data buffer descriptor are supported.

- **BUFFER ID**

Only buffer 0 will be supported.

- **BUFFER OFFSET**

Refers to the offset in the buffer to start reading data from. Should be less than the size of the buffersize, otherwise a check condition will be sent back with sense key set to ILLEGAL REQUEST and additional sense CODE set to INVALID FIELD IN CDB.

- **ALLOCATION LENGTH**

Refer to individual sections for the meaning of this term.

3.13.3 Data Returned

Data returned to the requestor depends on the MODE parameter.

3.13.3.1 Data Only Mode (02h)

In this mode, data is read from the device's sector buffer and returned to the requestor. Note that sector buffer in the SATA device is being used to emulate the SCSI READ BUFFER command, so the maximum length of data that can be written is 512 bytes. Valid fields in the CDB, apart from the MODE field, are BUFFER ID, BUFFER OFFSET (< 512) and ALLOCATION LENGTH (<= 512). A write buffer command may have been sent to the same buffer ID before it is read.

3.13.3.2 Descriptor mode (03h)

4 bytes of information will be returned to the requestor describing the requested buffer. These 4 bytes include the OFFSET BOUNDARY and the BUFFER CAPACITY. The BUFFER ID should be set to 0. For all other BUFFER ID's, all zeros will be returned. ALLOCATION LENGTH should be set to 4.

3.13.4 Miscellaneous Notes.

Sector buffer in a SATA device will be used to emulate the READ BUFFER command, so the size of the buffer is limited to 512 bytes for data buffer and echo buffers.

3.14 Write Buffer Command (3Bh)

The write buffer command is used with the read buffer command to determine the integrity of the target device's buffer memory and the physical interconnect that connects the target device and the initiator.

3.14.1 Command Summery

Table 3.14.1.15 Write Buffer command summery

Support Method	- Partially translated, partially emulated.		
ATA Opcodes	MODE	Supported Method	Translated ATA Opcode
	00h – Combined Header and Data Mode	Not Supported.	N/A
	01h – Vendor specific Mode	Not Supported.	N/A
	02h – Data Only mode	Translated to ATA Write Buffer Command	E8h
	04h–07h – Download micro code modes	Not Supported.	N/A
	0Ah – Echo Mode	Not Supported	N/A
	All other MODE values	Not Supported.	N/A

3.14.2 Read Buffer CDB format and supported fields

- **MODE**

Please refer Table 3.14.1.15 for a description of what MODEs are supported. Only data-buffer mode is supported.

- **BUFFER ID**

Only buffer 0 will be supported.

- **BUFFER OFFSET**

Refers to the offset in the buffer to start reading data from. Should be less than the size of the buffer size, otherwise a check condition will be sent back with sense key set to ILLEGAL REQUEST and additional sense CODE set to INVALID FIELD IN CDB.

- **PARAMETER LENGTH**

Refer to the maximum number of bytes transferred from the user buffer and should be less than the size of the buffer data is written to. If not, a check condition will be returned with sense key set to ILLEGAL REQUEST and additional sense code set to INVALID FIELD IN CDB.

3.14.3 Data written.

Data written to the device buffer depends on the MODE parameter.

3.14.3.1 Data Only Mode (02h)

In this mode, data is written to the device's sector buffer. Note that sector buffer in the SATA device is being used to emulate the SCSI WRITE BUFFER command, so the maximum length of data that can be written is 512 bytes. Valid fields in the CDB except the MODE field is 0 for BUFFER ID which should be 0, BUFFER OFFSET (< 512) and ALLOCATION LENGTH (<= 512).

3.14.4 Miscellaneous Notes.

Sector buffer in a SATA device will be used to emulate the WRITE BUFFER command, so the size of the buffer is limited to 512 bytes for data buffer and echo buffers.

3.15 Seek (06, 10) Command (0Bh, 2Bh)

The seek command has been made obsolete in ATA/ATAPI-7. All seek commands complete with success. No commands are sent to the ATA device.

3.15.1 Command Summary

Table 3.15.1.16 Seek command summary

Support Method	Command Obsolete in ATA/ATAPI-7, All all seek or extended seek commands complete with success
ATA Opcodes	None

3.15.2 Seek CDB format and supported fields

All fields are ignored.

3.15.3 Data Returned

None.

3.15.4 Miscellaneous Notes.

Could be supported with a READ VERIFY?

3.16 Rezero Unit (06) Command (01h)

The rezero SCSI command has been made obsolete in SBC 2 (January 12, 2004 Revision 12), and the seek command has been made obsolete in ATA/ATAPI-7. Return success to rezero commands and does not propagate the command to the target.

3.16.1 Command Summary

Table 3.16.1.17 Rezero command summary

Support Method	SCSI REZERO command too has been made obsolete since SBC-2, revision 12. The command completes with success
ATA Opcodes	None

3.16.2 Rezero CDB format and supported fields

Ignored.

3.16.3 Data Returned

None.

3.16.4 Miscellaneous Notes

Could be supported with a READ VERIFY?

3.17 Reserve Unit Command (16h)

The Reserve Unit, release unit SCSI commands have been made obsolete in SBC 2 (January 12, 2004 Revision 12). There is no meaningful translation of a SCSI reserve command to an ATA command. The translator shall return a check condition with sense key set to illegal request (05h) and additional sense code set to invalid command operation code (20h).

3.17.1 Command Summary

Table 3.17.1.18 Reserve Unit command summary

Support Method	Not supported.
ATA Opcodes	None

3.17.2 Reserve CDB format and supported fields

None, all parameters are ignored.

3.17.3 Data Returned

None.

3.17.4 Miscellaneous Notes

None.

3.17.5 SAS Affiliation Consideration

TBD.

3.18 Release Unit Command (17h) (06/10)

The release unit SCSI commands have been made obsolete in SBC 2 (January 12, 2004 Revision 12). There is no meaningful translation of a SCSI release command to an ATA command. The translator shall return a check condition with sense key set to illegal request (05h) and additional sense code set to invalid command operation code (20h).

3.18.1 Command Summary

Table 3.18.1.19 Reserve Unit/ Release Unit command summary

Support Method	Not supported.
ATA Opcodes	None

3.18.2 Reserve CDB format and supported fields

None, all parameters are ignored.

3.18.3 Data Returned

None.

3.18.4 Miscellaneous Notes

None.

3.18.5 SAS Affiliation Considerations

TBD.

3.19 Send Diagnostic Command (1Dh)

The send diagnostic command provides a mechanism for an application client to carry out diagnostic operations on the target, logical unit or both. Note that only drive targets can be serviced through this command. Also note that the complementary command, receive diagnostic command, is not supported.

3.19.1 Command Summary

Table 3.19.1.20 Send diagnostic command summary

Support Method	Default Self test – Translated to ATA Execute Device Diagnostic (90h). Background tests – Not Supported.
ATA Opcodes	90h

3.19.2 Send Diagnostic CDB format and supported fields

- UNITOFFL

Ignored.

- DEVOFFL

Ignored.

- SELFTEST

A value of 1 which indicates the default self test, be returned with a status of SUCCESS if the default self test completes without any errors, or a check condition with sense key set to hardware error if an error occurs.

- PF (Parameter format)

Ignored.

- PARAMETER LIST LENGTH

Ignored.

- SELF TEST CODE

If the SELFTEST bit is zero, SELF TEST CODE indicates which self-test to be carried out.

Table 3.19.2.21 Send diagnostic self-test code

<i>Value</i>	<i>Name of Test</i>	<i>Description of Test</i>	<i>Support</i>
000b	None.	Used when SELFTEST bit is set to 1.	Yes.
001b	Background short self test	Start a background self test. Results available in log page 10h.	No.
010b	Background Extended self test.	Start an extended self-test in background mode. Results available in Log page 10h.	No.
011b	Reserved.		
100b	Abort background self test.	If a self test is running in background mode, abort t.	No.
101b	Foreground short self test.	Start a self-test in foreground mode.	No.
110b	Foreground extended self test	Start an extended self-test in foreground mode.	No.

3.19.3 Data Returned

If SELFTEST bit set to 1, GOOD status will be returned if the default self test completes without any errors. Check condition will be returned if an error occurred.

3.19.4 Miscellaneous Notes

No mechanism will be provided to pass parameters from an application data-out buffer to the target.

Receive Diagnostic command is not supported.

3.20 Log Sense Command (4Dh)

The log sense command provides a mechanism, which an application can use to retrieve statistical or diagnostic results, or other operating data about a target or a logical unit. Note that only SMART data are returned through this command.

3.20.1 Command Summary

Table 3.20.1.22 Log Sense command summary

Support Method	For SMART data, i.e. Log Page 31h, translated to ATA SMART (B0h) command with feature register set to D0h.
ATA Opcodes	B0h

3.20.2 Log Sense CDB format and supported fields

- SP (Save Parameters)

This bit is ignored. Saving of parameters is not supported.

- PPC (Parameter Pointer Control)

This bit is ignored.

- PC (Page Control)

Table 3.20.2.23 Page Control Value

PC	LOG Sense Parameter value
00b	Current values.
01b	Cumulative Values.
10b	Default threshold values. N/A.
11b	Default cumulative values N/A.

- PAGE CODE

Table 3.20.2.24 Page Code Values

Page Code	Details
31h	The drive's SMART data with out the Page header field. Format corresponds to device's SMART data structure.
All other s	A check condition with sense key set to illegal request and additional sense code set to invalid field in CDB will be returned.

- PARAMETER POINTER

This filed is ignored.

- ALLOCATION LENGTH

Indicates the length of the use data-in buffer. Data up to or less the this size will be transferred.

3.20.3 Data Returned

SMART data log page data will be returned. If the page code is another value, a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB will be returned to the caller.

3.20.3.1 SMART data.

SMART data will be returned in the format supported by the device. No log page header will be appended to it. Valid PC values are 00b and 01b.

3.20.4 Miscellaneous Notes.

None.

3.21 Mode Sense (06)/(10) Command (1Ah)/(5Ah)

Mode sense command is used to query the device server about operational parameters of the target device or logical unit(s). The translator acts as the device server and emulates the SCSI behavior for SATA devices. Mode sense command is the complementary command of the MODE SELECT command. See section 4 for supported mode pages.

3.21.1 Command Summary

Table 3.21.1.25 Mode Sense (06)/(10) commands summary

Support Method	Mostly emulated. Some operational parameters in individual pages are gathered by issuing ATA commands. See section on Mode page for more details.
ATA Opcodes	No direct translations.

3.21.1.1 Mode Sense (06) CDB format and supported fields

- DBD

A value 0b indicates that one or more Block Descriptors can be returned in Mode Sense data. Currently only one Block descriptor is supported.

- PC (Page Control)

Table 3.21.1.26 Page Control (PC) values and their descriptions

PC Value	Type of Parameter	Supported
00b	Current Values	Yes
01b	Changeable Values	No
10b	Default Values	No
11b	Saved Values	No

- PAGE CODE

This field identifies the particular page requested. See section on Mode pages for further details.

- SUB PAGE CODE

This field identifies the sub page code within the page code specified by PAGE CODE field that is required by the client. See section on Mode pages for further details.

3.21.1.2 Mode Sense (10) CDB format and supported fields

- LLBAA (Long LBA Accepted)

- If this bit is set to 0, device server is prohibited from setting LONGLBA bit in the Mode parameter header. This indicates that mode parameter block descriptor is 8 bytes in length.
- If this bit is set to one, the device server is allowed to set LONGLBA bit in the Mode Parameter header. If LONGLBA bit is set to 1 when returning mode parameters it indicates that the block descriptors are 16 bytes in length.

3.21.2 Data Returned

The general format used to return data for a MODE SENSE command contains a Mode Parameter header, followed by one or more Block Descriptors, followed by a variable length mode page(s). Please refer SPC-3 for more details on the format of each of these sections. They are described briefly in SCSI-SATA translation context here.

3.21.2.1 Reply for Mode Sense (06) Command

3.21.2.1.1 Mode Sense Parameter Header (06)

- **MODE DATA LENGTH**
Indicates the number of bytes following this field that was transferred.
- **MEDIUM TYPE**
For Direct access device types, this field is set to 00h.
- **DEVICE SPECIFIC PARAMETER**
Refer SBC-2 for the format of the DEVICE SPECIFIC PARAMETER for block commands.
 - **DPOFUA**
 - If this bit is set to 0b, then it indicates that device server does not support DPO and FUA bits.
 - If this bit is set to 1b, then it indicates that the device server supports DPO and FUA bits.
 - **WP**
 - If this bit is set to 0b, it indicates that the device is write enabled. This is the only supported value for this bit.
 - If this bit is set to 1b, it indicates that the device is write protected, and this value is currently not supported.
- **BLOCK DESCRIPTOR LENGTH**
This field indicates the length of the block descriptor section. This value is obtained by multiplying the number of block descriptors by 8, for MODE SENSE (6) commands and MODE SENSE (10) commands, if LONGLBA bit is set to 0, or by 16 for MODE SENSE (10) commands and LONGLBA bit is set to 1.

3.21.2.1.2 Mode Sense Parameter Header (10)

Only the length of the MODE DATA LENGTH field and the BLOCK DESCRIPTOR LENGTH field are different for MODE SENSE (10) reply parameter header, except that LONGLBA bit is defined as below.

- **LONGLBA**
Describes the length of the block descriptors.
 - If set to 0b, block descriptors are 8 bytes long.
 - If set to 1b, block descriptors are 16 bytes long.

3.21.2.1.3 Mode Sense Block Descriptor (8 byte format)

- **NUMBER OF BLOCKS**

This field indicates the number of blocks being described by this block descriptor. Current implementation dictates that only one block descriptor is returned, so it will describe the blocks for the entire LUN. Therefore this field is equal to the Maximum LBA of the device minus one.

- **DENSITY CODE**

For direct access devices, this field is set to 00h.

- **BLOCK LENGTH**

Describes the block length for the section of the LUN described by this block descriptor. Since there is only one block descriptor, this describes the block length of the entire LUN. For direct access devices, this is set to 200h or 512 bytes per block.

3.21.2.1.4 Mode Sense Block Descriptor (16 byte format)

Only the length of the NUMBER OF BLOCKS field and the BLOCK LENGTH field changes between the 8-byte format and the 16-byte format, while the description of the field remains the same.

3.21.2.1.5 Mode Page format

Refer to the section on Mode Pages for the format of individual pages.

3.21.3 Miscellaneous Notes

Persistent saving of parameters is not supported. Only one block descriptor will be supported and will describe the settings for the entire logical unit.

3.22 Mode Select (06)/(10) Command (15h)/(55h)

Mode select command provides a mechanism for application clients to change the operating parameters of the target or a logical unit. This is the complementary command for MODE SENSE. The application client is advised to send a MODE SENSE command before issuing a MODE SELECT command for the same mode page, to determine the format, length changeable field etc. For all initiators, Logical Units of a target shall share the mode parameter header and the block descriptor and mode page values.

Currently Mode Page Policy VPD page is not implemented. Therefore Mode page policy is set to *'Shared'*, which indicates that only one copy of *'current'* values are maintained for all logical units of a target. After a logical unit reset, these values will reset to default values, since saved parameters are not supported. See section 4 for supported mode pages.

3.22.1 Command Summary

Table 3.22.1.27 Mode Select (06)/(10) commands summary

Support Method	Mostly emulated. Some operational parameters in individual pages are provided via ATA. See section on Mode page for more details.
ATA Opcodes	No direct translations.

3.22.1.1 Mode Select (06) CDB format and supported fields

- SP (Save parameters)

Save Parameter operation is not supported, there only a value of 0b is supported. If this bit is set to 1b, a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB, will be returned.

- PF (Page Format)

A page format of 0b indicates that modes pages are vendor specific, therefore a value of 0b is not supported. A value of 1b indicates that all page formats corresponds to SPC-3 and SBS-2 MODE PAGE formats. If this bit is set to 0b, a check condition with sense key set to illegal request and additional sense code set to invalid field in CDB, will be returned.

- PARAMETER LIST LENGTH

This field should be set to the length of the mode parameter list, transferred out of the client's data-out buffer.

3.22.1.2 Mode Select (10) CDB format and supported fields

Only the size of the PARAMETER LIST LENGTH field is different in the MODE SELECT (06) command, all other details remain the same.

3.22.2 Data Transferred

The general format used for data in MODE SELECT command contains a Mode Parameter header, followed by one or more Block Descriptors, followed by a variable length mode page(s). Please refer SPC-3 for more details on the format of each of these sections. They are described briefly in SCSI-SATA translation context here.

3.22.2.1 Format for Mode Select (06) Command

3.22.2.1.1 Mode Select Parameter Header (06)

- **MODE DATA LENGTH**

This field is reserved for MODE SELECT command and should be set to 00h.

- **MEDIUM TYPE**

For direct access devices, this field should be set to 00h.

- **DEVICE SPECIFIC PARAMETER**

This field is reserved for MODE SELECT command and should be set to 00h.

- **BLOCK DESCRIPTOR LENGTH**

This field indicates the length of the block descriptor section. This value is obtained by multiplying the number of block descriptors by 8, for MODE SELECT (6) commands and MODE SELECT (10) commands, if LONGLBA bit is set to 0, or by 16 for MODE SELECT (10) commands and LONGLBA bit is set to 1. Currently only one Block descriptor is supported, so this value should be set to 8 or 16.

3.22.2.1.2 Mode Select Parameter Header (10)

Only the length of the MODE DATA LENGTH field and the BLOCK DESCRIPTOR LENGTH field are different for MODE SELECT (10) reply parameter header, except that LONGLBA bit is defined as below.

- **LONGLBA** – Describes the length of the block descriptors.

- If set to 0b, block descriptors are 8 bytes long.
- If set to 1b, block descriptors are 16 bytes long.

3.22.2.1.3 Mode Select Block Descriptor (8 byte format)

- **NUMBER OF BLOCKS**

This field will be ignored as the capacity of the devices attached cannot be changed.

- **DENSITY CODE**

For direct access devices, this field should be set to 00h.

- **BLOCK LENGTH**

Describes the block length for the section of the LUN described by this block descriptor. Since there is only one block descriptor, this describes the block length of the entire LUN. For direct access devices, this is set to 200h or 512 bytes per block.

3.22.2.1.4 Mode Select Block Descriptor (16 byte format)

Only the length of the NUMBER OF BLOCKS field and the BLOCK LENGTH field changes between the 8-byte format and the 16-byte format, while the description of the field remains the same.

3.22.2.1.5 Mode Page format

Refer to the section on Mode Pages for the format of individual pages.

3.22.3 Miscellaneous Notes

Saving of parameters is not supported. Only one block descriptor will be supported and will describe the settings for the entire logical unit. Most MODE SELECT operations specified by the content of the

mode pages concerned will be turned to no ops due to the lack of equivalent commands or features in SATA. Refer to the section on Mode Pages for the details about the contents that are supported for each individual page.

4. SUPPORTED SCSI MODE PAGES

SCSI mode pages provide a mechanism to be used to set operating parameters for targets and logical units. MODE SENSE command is used to obtain operating parameters while MODE SELECT command is used to set operating parameters. In the context SCSI to SATA command translation most operating parameters defined by the contents of MODE PAGES are not supported due to lack of equivalent operations or features in SATA devices. The translator acts as the device server for all MODE SENSE and MODE SELECT commands and emulates the MODE PAGES listed in the table below.

In SCSI four types of mode page values are defined and they are described in Table 3.21.1.26. The '*save parameters*' operation is not supported. Default values are manufacturing time initial values and are not changeable. All pages are shared between all logical units in the target. Mode Page Policy VPD is not implemented; therefore sharing of Mode pages across logical units is implied.

Fields whose contents are changeable and ignored can be changed by an application client without an error being reported back.

Currently, most mode pages can be used for information purposes only. Use of the MODE SELECT command to change the parameters are turned into no ops. Where a behavioral change takes place, this specification will describe the ATA translation and the expected out come, where possible.

4.1 Commonly used SCSI Mode Pages

Table 4.1.1 lists the commonly used SCSI mode pages by operating system. This list is not exhaustive, and new entries will be added as deemed necessary.

Table 4.1.1 Commonly used SCSI mode pages

<i>Mode Page</i>	
<i>Page Code</i>	<i>Page Name</i>
01h	Read Write Error Recovery
02h	Disconnect-Reconnect
03h	Format Device
04h	Rigid Disk Geometry
08h	Caching
0Ah	Control
1Ch	Informational Exception Control
3Fh	Return All Pages

4.2 Read-Write Error Recovery Mode Page (01h)

This page contains the parameters used by the device server during error condition when read write commands are executed. Read write operations include all variants of the READ command, the WRITE command and WRITE AND VERIFY commands. Note that an ignored bit can be changed, unless otherwise stated, but the change will have no effect on the operation of the target.

4.2.1 Read-Write Error Recovery Mode Page format and supported fields

- PS (Parameters Savable)

This bit will be set to 0b to indicate that parameters cannot be save in a non-volatile location. It is reserved for MODE SELECT commands.
- DCR (Disable Correction)

This bit will be set to 0b and ignored as data correction is done at drive level.
- DTE (Disable Transfer on Error0)

This bit is ignored and set to 0 to indicate that data transfers will not be terminated at the detection of a recovered error.
- PER (Post Error)

This bit is ignored and set to 0 to indicate that recovered errors not reported to the user.
- EER (Enable Early Recovery)

This bit is ignored set to 0 to indicate that the device server shall use the error recovery procedure that minimizes the risk of mis-detection or mis-correction.
- RC (Read Continuous)

This bit is ignored and set to 0 to indicate that error recovery delays are acceptable and that Data will not be fabricated.
- TB (Transfer Block)

This bit is ignored and set to 0 to indicate that a logical block that is not recovered within the recovery limits specified shall not be transferred to the application client.
- ARRE (Automatic Read Reallocation)

This bit is ignored and will be set to 1 to indicate that automatic read reallocation of defective blocks will always be enabled at drive level.
- AWRE (Automatic Write Reallocation)

This bit is ignored and will be set to 1 to indicate that automatic Write reallocation of defective blocks will always be enabled at drive level.
- READ RETRY COUNT

This field will be set to 00h and ignored. This field identifies the maximum number of times the operation will be retried by the device server in its recovery algorithm.
- WRITE RETRY COUNT

This field will be set to 00h and ignored. This field identifies the maximum number of times the operation will be retried by the device server in its recovery algorithm.
- RECOVERY TIME LIMIT

This field is ignored and set to 00h. This field specifies in increments of one millisecond, the maximum duration the device server shall use for error recovery.

4.3 Format Mode Page (03h)

(Made obsolete in SPI-3, Do we want this?)

- TRACKS PER ZONE

This value would be set to 128. It specifies the number of tracks per zone in dividing the capacity of the drive for the purpose of allocating alternate sectors.

- ALTERNATE SECTORS PER ZONE

This field will be set to 00h and is ignored. This specifies the number of sectors the device server shall reserve per zone for defect handling.

- ALTERNATE TRACKS PER ZONE

This field will be set to 00h and is ignored. This specifies the number of the tracks device server shall reserve per zone for defect handling.

- ALTERNATE TRACKS PER LOGICAL UNIT

This field will be set to 00h and is ignored. This specifies the number of the tracks device server shall reserve per logical units for defect handling.

- SECTORS PER TRACK

This value will be set to 128. Self-explanatory.

- BLOCK SIZE

This value will be set to 512 bytes. Self-explanatory.

- INTERLEAVE

This value will be set to 00h. This specifies the same value from the last format operation. Only a value of 00h is supported in the original format command, which implies that the default interleave mechanism for the drive will be used.

- TRACK SKEW FACTOR

This value will be set to 01h. specifies the number of physical sectors between the last logical block of one cylinder and the first logical block the next sequential cylinder.

- SSEC (Soft Sector Formatting)

Set to 0 and ignored. Mutually exclusive with HSEC bit.

- HSEC (Hard Sector formatting)

This bit will be set to 1b to indicate that hard sector formatting will be used.

- RMB (Removable media)

This bit will be set to 0 and ignored.

- SURF (Surface)

This bit will be set to 0 and ignored.

4.4 Rigid Disk Device Geometry Page (04h)

- NUMBER OF CYLINDERS
This field will be set to the number of physical cylinders used for data storage.
- NUMBER OF HEADS
This field will be set to the number of heads used for data storage.
- START CYLINDER FOR WRITE PRECOMPENSATION
This value will be set to 0 and ignored.
- START CYLINDER REDUCED WRITE CURRENT
This value will be set to 0 and ignored.
- DEVICE STEP RATE
This value will be set to 0 and ignored.
- LANDING ZONE CYLINDER
This value will be set to 0 and ignored.
- ROTATIONAL OFFSET
This value will be set to 0 and ignored.
- MEDIUM ROTATION RATE
This value will be set to the rotation rate of the device. Units are rotation per minute.

4.5 Caching Mode Page (08h)

Caching Mode page (08h) defines parameters that affect the behavior of the device cache and the caching policy used.

- PS (Parameters Savable)
This bit will be set to 0b to indicate that parameters cannot be save in a non-volatile location. It is reserved for MODE SELECT commands.
- RCD (Read Cache Disable)
This bit will be set to 0 and ignored. Set DRA bit to prevent read ahead, which is the equivalent of enabling and disabling of read cache in SATA.
- MF (Multiplication Factor)
This bit is set to 0 and is ignored.
- WCE (Write Cache Enable)
 - By default, this bit will be set to 1 to indicate that write caching is enabled.
 - If this bit is set to 0 using a MODE SELECT command, it will be translated into a SET FEATURES command 82h to disable write cache.
 - If this bit is set to 1 using a MODE SELECT command, it will be translated into a SET FEATURES command 02h to enable write cache.
 - This bit it set to the default values after a logical unit reset or a target reset.

- SIZE

This bit is set to 0 and is ignored. This bit indicates whether CACHE SEGMENT SIZE field or NUMBER OF CACHE SEGMENTS field to use for caching algorithms.
- DISC (Discontinuity)

This bit is set to 0 and is ignored. This bit defines whether prefetches can be continued across discontinuities.
- CAP (Caching analysis permitted)

This bit is set to 0 to indicate that caching analysis is disabled and ignored for a MODE SELECT command.
- ABPF (Abort Prefetch)

This bit is set to 0 and ignored. This bit is used in conjunction with the DRA bit and defines how prefetch operation is controlled when a new command is received.
- IC (initiator Control)

This bit is set to 0 and ignored. Initiator control bit set to 1 specifies that the device server use the NUMBER OF CACHE SEGMENTS or the CACHE SEGMENT SIZE, depending on the SIZE bit to control the caching algorithm.
- WRITE RETENTION PRIORITY

This field is set to zero and ignored.
- DEMAND READ RETENTION PRIORITY

This field is set to zero and ignored.
- DISABLE PRE-FETCH TRANSFER LENGTH

This field is set to zero and ignored. This field specifies the cut off transfer length in number blocks to disable prefetch for long transfers.
- MINIMUM PRE-FETCH

This field is set to zero and ignored.
- MAXIMUM PRE-FETCH

This field is set to zero and ignored.
- MAXIMUM PRE-FETCH CEILING

This field is set to zero and ignored.
- DRA (Disable Read Ahead)
 - By default, this bit will be set to 0 to indicate that read look-ahead is enabled.
 - If this bit is set to 1 using a MODE SELECT command, it will be translated into a SET FEATURES command 55h to disable read look-ahead.
 - If this bit is set to 0 using a MODE SELECT command, it will be translated into a SET FEATURES command AAh to enable read look-ahead.
 - This bit is set to the default values after a logical unit reset or a target reset.
- LBCSS (Logical block cache segment size)

This field is set to zero and ignored. This bit specifies the units to be used for CACHE SEGMENT SIZE field.

- FSW (Force Sequential Writes)

This field is set to zero and ignored.
- NUMBER OF CACHE SEGMENTS

This field is set to zero and ignored.
- CACHE SEGMENT SIZE

This field is set to zero and ignored.
- NON-CACHE SEGMENT SIZE

This field is set to zero and ignored.

4.6 Informational Exceptions Control Mode Page (1Ch)

This page allows the initiator to specify the methods used to control the reporting and operation under specific informational exception conditions. This page applies to informational exceptions that report an additional sense code of FAILURE PREDICTION THRESHOLD EXCEEDED or WARNING to the application client.

- PS (Parameters Savable)

This bit will be set to 0b to indicate that parameters cannot be save in a non-volatile location. It is reserved for MODE SELECT commands.
- LOGERR (Log Error)

This bit is set to 0 and ignored. Indicates that logging of informational exceptions is vendor specific.
- TEST

This bit is set to 0b to indicate reporting false device failure notifications is not supported.
- DEXCPT (Disable Exception Control)

Set to 0, enables informational exceptions reporting using the method defined by MRIE. This is the default value. SMART is enabled if this bit is 0b.

Set to 1, disables informational exceptions reporting using the method defined by MRIE. SMART is disabled if this bit is 0b.
- EWASC (enable Warning)

This bit is set to 0 and is ignored. Reporting of warning will be disabled.
- EBF (Enable Background Functions)

This bit is set to 0 and is ignored. Enabling of background functions will be disabled. PERF
- (Performance)

This bit is set to 0 and is ignored.
- MIRE (Method of reporting Informational Exceptions)

Only a value of 6h (report on request) is supported.
- INTERVAL TIMER

This bit is set to 0 and is ignored. This field defines the interval in 100 milliseconds that the device will be polled for SMART threshold conditions. Conditions will be reported after the interval has elapsed. Check section 6 for more details on SMART.

- REPORT COUNT

This field is set to 0 and ignored. The number of times a condition can be reported is unlimited.

5. SCSI TASK MANAGEMENT TRANSLATION

This section describes the translation of SCSI task management message to SATA equivalents. Note that due to architectural differences, not all task management messages can be successfully translated to ATA commands or control operations. Please refer SPI-4 specification for specific details of task management operations.

5.1 Abort Task

'Abort task' task management operation is not supported for SATA devices. A status of FAILED will be returned for the abort task request. Previously established conditions such CA, Mode Parameters, reservations shall not change. The particular task in consideration may complete as any other request, with or without a successful completion status.

5.2 Abort Task Set

'Abort task set' operation is not supported for SATA devices. A status of FAILED will be returned for the 'abort task set' request. Previously established conditions such, Mode Parameters, reservations shall not change. The particular tasks in consideration may complete as any other request, with or without a successful completion status.

5.3 Logical Unit Reset

'Logical Unit Reset' operation is not supported for SATA devices. A status of FAILED will be returned for the 'Logical Unit Reset' request. Previously established conditions such, Mode Parameters, reservations shall not change. Tasks submitted for this logical unit may complete as any other request, with or without a successful completion status.

5.4 Target Reset

Target reset operation is supported by the following sequence.

- All tasks in the task set for the target will be aborted.
- Any CA condition will be cleared.
- A device reset will be carried out for the specified target by performing a link reset.
- The device will be brought to its operational mode and all Mode Parameters will be set to default values.

5.5 Bus Reset

Bus reset operation is supported performing a link reset of all associated targets on the logical bus.

5.6 Clear Task Set

'Clear task set' operation is not supported for SATA devices. A status of FAILED will be returned for the 'clear task set' request. Previously established conditions such, Mode Parameters, reservations shall not change. The particular tasks in consideration may complete as any other request, with or without a successful completion status.

6. SMART HANDLING

This section describes the SMART features used by most drive targets to monitor the health of the drives by application software.

Needs work. TBD

6.1 General Description of SMART

Provide a method for an application client to query for SMART attribute threshold exceeding conditions. The host will be notified in a manner specified by the Informational Exception Mode page parameters. This page is also used to enable or disable SMART feature in full. Once the application software becomes aware of a threshold exceeding condition, it is the duty of application software to retrieve SMART data, analyze it and take necessary actions. A simple method is provided to retrieve SMART data. Following sub sections address these aspects in more details.

6.2 Host Notification

The method used to notify the host of a SMART threshold exceeding condition is specified in the Informational Exceptions Mode Page, using the MRIE field. Only a value of 06h is supported for this field as described in section 4.6. A value of 06h indicates that the application client will send unsolicited request sense commands described in 3.6, to monitor SMART errors. When an unsolicited request sense command is recieved, issue an ATA SMART RETURN STATUS command if the target device is SMART capable. (ATA opcode B0h with features register set to DAh). Upon receiving results, if a threshold exceeding condition has been detected and met, sense data with the sense key set to no sense and additional sense code set to general hardware failure will be returned. (ASC 5Dh, ASCQ 10h).

6.3 Disabling, enabling SMART

To enable or disable SMART reporting use the DEXCPT bit in the Informational Exceptions Mode Page as described in section 4.6, using a mode select command. A value of 0, which is the default, indicates to enable informational exceptions reporting while a value of 1 indicates to disable informational exceptions reporting. Do not issue SMART RETURN STATUS command, if DEXCPT bit is set to 1.

6.4 Retrieving SMART data from targets

Application software should use LOG SENSE command described in section 3.20, to retrieve 512 byte SMART unaltered data from drives. Only page 31h will be supported under LOG SENSE command and only current values as specified by the PC field (00b) will be reported to the client. The page header will not be appended to SMART data. Also the LOG SELECT command is not supported. Transfer length can be controlled using the ALLOCATION LENGTH field in the LOG SENSE CDB.